

## EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S6	2852	database\$1 and reconcil\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:14
S7	413	database\$1 with reconcil\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:20
S8	324	(database\$1 with reconcil\$3) and (log\$4 or history)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:10
S9	71	(database\$1 with reconcil\$3) and ((search\$3 or quer\$3) with (log\$4 or history))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:28
S10	57	((database\$1 with reconcil\$3) and ((search\$3 or quer\$3) with (log\$4 or history))) and (increment\$3 or increas\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:15
S11	23	((database\$1 with reconcil\$3) and synchroniz\$4) and ((search\$3 or quer\$3) with (log\$1 or history))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:15
S12	15	(((database\$1 with reconcil\$3) and synchroniz\$4) and ((search\$3 or quer\$3) with (log\$1 or history))) and increment\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:15
S13	141	(database\$1 with reconcil\$3) and synchroniz\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:16
S14	54	(database\$1 with reconcil\$3) and synchroniz\$4 and increment\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:28
S15	2922	database\$1 with (backup\$3 or (back adj up\$4))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:24
S16	30594	database\$1 with updat\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:27
S17	292	database\$1 with updat\$3 with consistent	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:27

## EAST Search History

S18	145	(database\$1 with updat\$3 with consistent) and increment\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:28
S19	102	((database\$1 with updat\$3 with consistent) and increment\$3) and (log\$1 or history)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/01 15:28
S20	309	(synchroniz\$5 and database\$1).ti.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 06:02
S21	1201	(synchroniz\$5 and database\$1).ab.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 13:25
S22	268	((synchroniz\$5 and database\$1).ab.) and (sequenc\$3 or sequential)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 08:28
S23	189	((synchroniz\$5 and database\$1).ab.) and (sequenc\$3 or sequential) ) and (increment\$3 or increas\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 06:04
S24	111	((synchroniz\$5 and database\$1).ab.) and (record\$1 with (ID or identifier\$1))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 08:26
S25	1207	help adj desk\$1	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 08:21
S26	0	((synchroniz\$5 and database\$1).ti.) and (help adj desk\$1)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 08:22
S27	5	((synchroniz\$5 and database\$1).ab.) and (help adj desk\$1)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 08:22
S28	76	((synchroniz\$5 and database\$1).ab.) and (record\$1 with (ID or identifier\$1)) and "707"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 08:29
S29	55	((synchroniz\$5 and database\$1).ab.) and (record\$1 with (ID or identifier\$1))) and (sequenc\$3 or sequential)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 08:29

## EAST Search History

S30	40	((((synchroniz\$5 and database\$1).ab.) and (record\$1 with (ID or identifier\$1))) and (sequenc\$3 or sequential) ) and "707"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 08:29
S31	30	(synchroniz\$5 and database\$1).ab. and ((increas\$3 or increment\$3) near2 (one or "6516327".PN.))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2004/06/02 13:26
S34	388	non-confidential	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/02/16 07:55
S36	3827	(remote adj database)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/09/16 08:10
S37	1368	(sequence adj number) adj field\$1	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/09/16 08:11
S38	3	(problem adj identifier) adj field\$1	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/09/16 08:21

**PORTAL** Reports (Full Service) Reports (Limited Service, Free) Log In

Search:  The Guide

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Published before January 2002 Found 32 of 124,448

Terms used [synchroniz](#) [database](#) [sequence number](#) [problem](#) [record](#) [greater](#) [append](#)

Sort results by   [Save results to a Binder](#) Try an [Advanced Search](#)  
 Display results   [Search Tips](#) Try this search in [The ACM Guide](#)

[Open results in a new window](#)

Results 1 - 20 of 32 Result page: [1](#) [2](#) [next](#)

Relevance scale

**1 Reliable scheduling in a TMR database system**

Frank M. Pittelli, Hector Garcia-Molina  
January 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 1

Publisher: ACM Press Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#), [review](#)

Full text available: [pdf\(2.14 MB\)](#)

A Triple Modular Redundant (TMR) system achieves high reliability by replicating data and all processing at three independent nodes. When TMR is used for database processing all nonfaulty computers must execute the same sequence of transactions, and this is ensured by a collection of processes known as schedulers. In this paper we study the implementation of efficient schedulers through analysis of various enhancements such as null transactions and message batching. The sch ...

**2 Fast detection of communication patterns in distributed executions**

Thomas Kunz, Michiel F. H. Seuren  
November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Publisher: IBM Press Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Full text available: [pdf\(4.21 MB\)](#)

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

**3 Query evaluation techniques for large databases**

Goetz Graefe  
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Publisher: ACM Press Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#), [review](#)

Full text available: [pdf\(9.37 MB\)](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

**4 ARIES/CSA: a method for database recovery in client-server architectures**

C. Mohan, Inderpal Narang  
May 1994 **ACM SIGMOD Record, Proceedings of the 1994 ACM SIGMOD international conference on Management of data SIGMOD '94**, Volume 23 Issue 2

Publisher: ACM Press Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

Full text available: [pdf\(1.33 MB\)](#)

This paper presents an algorithm, called ARIES/CSA (Algorithm for Recovery and Isolation Exploiting

Semantics for Client-Server Architectures), for performing recovery correctly in client-server (CS) architectures. In CS, the server manages the disk version of the database. The clients, after obtaining database pages from the server, cache them in their buffer pools. Clients perform their updates on the cached pages and produce log records. The log records are buffered loca ...

**5 Distributed systems - programming and management: On remote procedure call**

Patricia Gomes Soares

November 1992 **Proceedings of the 1992 conference of the Centre for Advanced Studies on Collaborative research - Volume 2**

Publisher: IBM Press

Full text available:  pdf(4.52 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)



The Remote Procedure Call (RPC) paradigm is reviewed. The concept is described, along with the backbone structure of the mechanisms that support it. An overview of works in supporting these mechanisms is discussed. Extensions to the paradigm that have been proposed to enlarge its suitability, are studied. The main contributions of this paper are a standard view and classification of RPC mechanisms according to different perspectives, and a snapshot of the paradigm in use today and of goals for t ...

**6 System R: relational approach to database management**



M. M. Astrahan, M. W. Blasgen, D. D. Chamberlin, K. P. Eswaran, J. N. Gray, P. P. Griffiths, W. F. King, R. A. Lorie, P. R. McJones, J. W. Melhi, G. R. Putzolu, I. L. Traiger, B. W. Wade, V. Watson

June 1976 **ACM Transactions on Database Systems (TODS)**, Volume 1 Issue 2

Publisher: ACM Press

Full text available:  pdf(3.18 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



System R is a database management system which provides a high level relational data interface. The system provides a high level of data independence by isolating the end user as much as possible from underlying storage structures. The system permits definition of a variety of relational views on common underlying data. Data control features are provided, including authorization, integrity assertions, triggered transactions, a logging and recovery subsystem, and facilities for maintaining ...

**Keywords:** authorization, data structures, database, index structures, locking, nonprocedural language, recovery, relational model

**7 ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging**



C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1

Publisher: ACM Press

Full text available:  pdf(5.23 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)



DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

**Keywords:** buffer management, latching, locking, space management, write-ahead logging

**8 A Majority consensus approach to concurrency control for multiple copy databases**



Robert H. Thomas

June 1979

**ACM Transactions on Database Systems (TODS)**, Volume 4 Issue 2

Publisher: ACM Press

Full text available:  pdf(2.32 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



A "majority consensus" algorithm which represents a new solution to the update synchronization problem for multiple copy databases is presented. The algorithm embodies distributed control and can function effectively in the presence of communication and database site outages. The correctness of the algorithm is demonstrated and the cost of using it is analyzed. Several examples that illustrate aspects of the algorithm operation are included in the Appendix.

**Keywords:** clock synchronization, computer networks, concurrency control, distributed computation, distributed control, distributed databases, multiprocess systems, update synchronization

**9 Remote pipes and procedures for efficient distributed communication**



David K. Gifford, Nathan Glasser

August 1988

**ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 3

Publisher: ACM Press

Full text available:  pdf(2.08 MB)

Additional Information: full citation, abstract, references, citations, index terms, review

We describe a new communication model for distributed systems that combines the advantages of remote procedure call with the efficient transfer of bulk data. Three ideas form the basis of this model. First, remote procedures are first-class values which can be freely exchanged among nodes, thus enabling a greater variety of protocols to be directly implemented in a remote procedure call framework. Second, a new type of abstract object, called a pipe, allows bulk data and in ...

10 UIO: a uniform I/O system interface for distributed systems

David R. Cheriton

January 1987

**ACM Transactions on Computer Systems (TOCS)**, Volume 5 Issue 1

Publisher: ACM Press

Full text available:  pdf(3.20 MB)

Additional Information: full citation, abstract, references, citations, index terms, review

A uniform I/O interface allows programs to be written relatively independently of specific I/O services and yet work with a wide variety of the I/O services available in a distributed environment. Ideally, the interface provides this uniform access without excessive complexity in the interface or loss of performance. However, a uniform interface does not arise from careful design of individual system interfaces alone; it requires explicit definition. In this paper, the UIO (unifo ...

11 Crash recovery in client-server EXODUS

Michael J. Franklin

June 1992

**ACM SIGMOD Record, Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.50 MB)

Additional Information: full citation, abstract, references, citations, index terms, review

In this paper, we address the correctness and performance issues that arise when implementing logging and crash recovery in a page-server environment. The issues result from two characteristics of page-server systems: 1) the fact that data is modified and cached in client database buffers that are not accessible by the server, and 2) the performance and cost trade-offs that are inherent in a client-server environment. We describe a recovery system that we have implemented for the client-server ...

12 File servers for network-based distributed systems

Liba Svobodova

December 1984

**ACM Computing Surveys (CSUR)**, Volume 16 Issue 4

Publisher: ACM Press

Full text available:  pdf(4.23 MB)

Additional Information: full citation, references, citations, index terms, review

13 An overview of the SR language and implementation

Gregory R. Andrews

Titus Purdin

January 1988

**ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 10

Issue 1

Publisher: ACM Press

Full text available:  pdf(2.98 MB)

Additional Information: full citation, abstract, references, citations, index terms, review



SR is a language for programming distributed systems ranging from operating systems to application programs. On the basis of our experience with the initial version, the language has evolved considerably. In this paper we describe the current version of SR and give an overview of its implementation. The main language constructs are still resources and operations. Resources encapsulate processes and variables that they share; operations provide the primary mechanism for process interaction. ...

14 Replicated data management in distributed database systems

Sang Hyuk Son

November 1988

**ACM SIGMOD Record**, Volume 17 Issue 4

Publisher: ACM Press

Full text available:  pdf(835.25 KB)

Additional Information: full citation, abstract, references, citations, index terms, review



Replication is the key factor in improving the availability of data in distributed systems. Replicated data is stored at multiple sites so that it can be accessed by the user even when some of the copies are not available due to site failures. A major restriction to using replication is that replicated copies must behave like a single copy, i.e., mutual consistency as well as internal consistency must be preserved. Synchronization techniques for replicated data in distributed database syste ...

15 The design of POSTGRES

Michael Stonebraker, Lawrence A. Rowe

June 1986

**ACM SIGMOD Record, Proceedings of the 1986 ACM SIGMOD international conference on Management of data SIGMOD '86**, Volume 15 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.91 MB)

Additional Information: full citation, abstract, references, citations, index terms

This paper presents the preliminary design of a new database management system, called POSTGRES, that is the successor to the INGRES relational database system. The main design goals of the new system are to provide better support for complex objects, provide user extensibility for data types, operators and access methods, provide facilities for active databases (i.e., alerters and triggers) and inferencing including forward- ...

16 A recovery algorithm for a high-performance memory-resident database system

Tobin J. Lehman, Michael J. Carey

December 1987

**ACM SIGMOD Record , Proceedings of the 1987 ACM SIGMOD international conference on Management of data SIGMOD '87**, Volume 16 Issue 3

Publisher: ACM Press

Full text available:  pdf(1.50 MB)

Additional Information: full citation, abstract, references, citations, index terms

With memory prices dropping and memory sizes increasing accordingly, a number of researchers are addressing the problem of designing high-performance database systems for managing memory-resident data. In this paper we address the recovery problem in the context of such a system. We argue that existing database recovery schemes fall short of meeting the requirements of such a system, and we present a new recovery mechanism which is designed to overcome their shortcomings. The proposed mecha ...

17 Real-time protocol analysis for detecting link-state routing protocol attacks

Ho-Yen Chang, S. Felix Wu, Y. Frank Jou

February 2001

**ACM Transactions on Information and System Security (TISSEC)**, Volume 4 Issue 1

Publisher: ACM Press

Full text available:  pdf(252.10 KB)

Additional Information: full citation, abstract, references, citations, index terms

A real-time knowledge-based network intrusion-detection model for a link-state routing protocol is presented for the OSPF protocol. This model includes three layers: a data process layer to parse packets and dispatch data; and event abstractor to abstract predefined real-time events for the link-state routing protocol; and an extended timed finite state machine to express the real-time behavior of the protocol engine and to ...

**Keywords:** OSPF attacks, event correlation, knowledge-based IDS, link-state routing protocol security, real-time misuse intrusion detection, real-time network protocol analysis, timed finite state machine

18 Optimistic recovery in distributed systems

Rob Strom, Shaula Yemini

August 1985

**ACM Transactions on Computer Systems (TOCS)**, Volume 3 Issue 3

Publisher: ACM Press

Full text available:  pdf(1.75 MB)

Additional Information: full citation, abstract, references, citations, index terms, review

*Optimistic Recovery* is a new technique supporting application-independent transparent recovery from processor failures in distributed systems. In optimistic recovery communication, computation and checkpointing proceed asynchronously. Synchronization is replaced by *causal dependency tracking*, which enables a posteriori reconstruction of a consistent distributed system state following a failure using *process rollback* and *message replay*.

Because there is no s ...

19 Integrating security in a large distributed system

M. Satyanarayanan

August 1989

**ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 3

Publisher: ACM Press

Full text available:  pdf(2.90 MB)

Additional Information: full citation, abstract, references, citations, index terms, review

Andrew is a distributed computing environment that is a synthesis of the personal computing and timesharing paradigms. When mature, it is expected to encompass over 5,000 workstations spanning the Carnegie Mellon University campus. This paper examines the security issues that arise in such an environment and describes the mechanisms that have been developed to address them. These mechanisms include the logical and physical separation of servers and clients, support for secure communication ...

20 Implementation of Argus

B. Liskov, D. Curtis, P. Johnson, R. Scheifer

November 1987

**ACM SIGOPS Operating Systems Review , Proceedings of the eleventh ACM Symposium on Operating systems principles SOSP '87**, Volume 21 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.34 MB)

Additional Information: full citation, abstract, references, citations, index terms

Argus is a programming language and system developed to support the construction and execution of distributed programs. This paper describes the implementation of Argus, with particular emphasis on the way we implement atomic actions, because this is where Argus differs most from other implemented systems. The paper also discusses the performance of Argus. The cost of actions is quite reasonable, indicating that action systems like Argus are practical.

Results 1 - 20 of 32

Result page: [1](#) [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

Published before January 2002  
Terms used [asynchronous](#) [database](#) [sequence number](#) [problem](#) [record](#) [greater](#) [append](#)

Sort results by [relevance](#) [expanded form](#)

Display results [Save results to a Binder](#) [Search Tips](#) [Open results in a new window](#)

Try an [Advanced Search](#)  
Try this search in [The ACM Guide](#)

Results 21 - 32 of 32

Result page: [previous](#) [1](#) [2](#)

Relevance scale

**21 Ordered and reliable multicast communication**  
Hector Garcia-Molina, AnneMarie Spauster  
August 1991 **ACM Transactions on Computer Systems (TOCS)**, Volume 9 Issue 3  
Publisher: ACM Press  
Full text available:  [pdf\(1.90 MB\)](#) Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

**Keywords:** message ordering, multicast communication

**22 A coherent distributed file cache with directory write-behind**  
Timothy Mann, Andrew Birrell, Andy Hisgen, Charles Jerian, Garret Swart  
May 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 2  
Publisher: ACM Press  
Full text available:  [pdf\(3.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#), [review](#)

Extensive caching is a key feature of the Echo distributed file system. Echo client machines maintain coherent caches of file and directory data and properties, with write-behind (delayed write-back) of all cached information. Echo specifies ordering constraints on this write-behind, enabling applications to store and maintain consistent data structures in the file system even when crashes or network faults prevent some writes from being completed. In this paper we describe ...

**Keywords:** coherence, file caching, write-behind

**23 Log files: an extended file service exploiting write-once storage**  
R. Finlayson, D. Cheriton  
November 1987 **ACM SIGOPS Operating Systems Review, Proceedings of the eleventh ACM Symposium on Operating systems principles SOSP '87**, Volume 21 Issue 5  
Publisher: ACM Press  
Full text available:  [pdf\(1.07 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

A log service provides efficient storage and retrieval of data that is written sequentially (append-only) and not subsequently modified. Application programs and subsystems use log services for recovery, to record security audit trails, and for performance monitoring. Ideally, a log service should accommodate very large, long-lived logs, and provide efficient retrieval and low space overhead. In this paper, we describe the design and implementation of the Clio log service. Clio pr ...

**24 Network Protocols**  
Andrew S. Tanenbaum  
December 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 4  
Publisher: ACM Press  
Full text available:  [pdf\(3.37 MB\)](#) Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

**25 Proceedings of the SIGNUM conference on the programming environment for development of numerical software**  
March 1979 **ACM SIGNUM Newsletter**, Volume 14 Issue 1  
Publisher: ACM Press

Full text available: [pdf\(5.02 MB\)](#)

Additional Information: [full citation](#)

**26 The Zebra striped network file system**



John H. Hartman, John K. Ousterhout

August 1995

**ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 3

Publisher: ACM Press

Full text available: [pdf\(2.76 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#), [review](#)

Zebra is a network file system that increases throughput by striping the file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this ...

**Keywords:** RAID, log-based striping, log-structured file system, parity computation

**27 The Zebra striped network file system**



John H. Hartman, John K. Ousterhout

December 1993

**ACM SIGOPS Operating Systems Review, Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available: [pdf\(1.93 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

Zebra is a network file system that increases throughput by striping file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this increase ...

**28 Implementing sequentially consistent shared objects using broadcast and point-to-point communication**



Alan Fekete, M. Frans Kaashoek, Nancy Lynch

January 1998

**Journal of the ACM (JACM)**, Volume 45 Issue 1

Publisher: ACM Press

Full text available: [pdf\(257.13 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

This paper presents and proves correct a distributed algorithm that implements a sequentially consistent collection of shared read/update objects. This algorithm is a generalization of one used in the Orca shared object system. The algorithm caches objects in the local memory of processors according to application needs; each read operation accesses a single copy of the object, while each update accesses all copies. The algorithm uses broadcast communication when it sends messages to replic ...

**Keywords:** Orca programming language, distributed shared memory, formal methods, input/output automata, ordered multicast, replicated data

**29 Frangipani: a scalable distributed file system**



Chandramohan A. Thekkath, Timothy Mann, Edward K. Lee

October 1997

**ACM SIGOPS Operating Systems Review, Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: [pdf\(2.20 MB\)](#)

Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

**30 The Xpress transfer protocol (XTP)—a tutorial**



Robert M. Sanders, Alfred C. Weaver

October 1990

**ACM SIGCOMM Computer Communication Review**, Volume 20 Issue 5

Publisher: ACM Press

Full text available: [pdf\(1.18 MB\)](#)

Additional Information: [full citation](#), [citing](#), [index terms](#)

**31 Secure agreement protocols: reliable and atomic group multicast in rampart**



Michael K. Reiter

November 1994

**Proceedings of the 2nd ACM Conference on Computer and communications security**

Publisher: ACM Press

Full text available: [pdf\(1.48 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

Reliable and atomic group multicast have been proposed as fundamental communication paradigms to support secure distributed computing in systems in which processes may behave maliciously. These protocols enable messages to be multicast to a group of processes, while ensuring that all honest group members deliver the same messages and, in the case of atomic multicast, deliver these messages in the same order. We present new reliable and atomic group multicast protocols for asynchronous distr ...

 **32 IP paging service for mobile hosts**

R. Ramjee, L. Li, T. La Porta, S. Kasera  
July 2001

**Proceedings of the 7th annual international conference on Mobile computing and networking**

Publisher: ACM Press

Full text available:  pdf(355.04 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In wireless networks, mobile hosts must update the network with their current location in order to get packets delivered. Paging facilitates efficient power management at the mobile host by allowing the host to update the network less frequently at the cost of providing the network with only approximate location information. The network determines the exact location of a mobile host through paging before delivering packets destined to the mobile host. In this paper, we propose the concept of p ...

Results 21 - 32 of 32

Result page: [previous](#) **1** [2](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

**Search Results**[BROWSE](#)[SEARCH](#)[IEEE Xplore GUIDE](#)

Results for "( ( synchroniz\* database &lt;in&gt;metadata ) ) &lt;and&gt; (pyr &gt;= 1950 &lt;and&gt; pyr &lt;= 20...)"

[e-mail](#)

Your search matched 2 of 1328352 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending** order.» [Search Options](#)[View Session History](#)[Modify Search](#)[New Search](#)[Search](#) Check to search only within this results set» **Key**Display Format:  Citation  Citation & Abstract**IEEE JNL** IEEE Journal or Magazine[view selected items](#) [Select All](#) [Deselect All](#)**IEE JNL** IEE Journal or Magazine

- 1. **Design and analysis of communication network for distributed SCADA systems in electric railways]**  
Qian Wang; Qingquan Qian;  
[Power Engineering Society Winter Meeting, 2000. IEEE](#)  
Volume 3, 23-27 Jan. 2000 Page(s):2062 - 2065 vol.3  
Digital Object Identifier 10.1109/PESW.2000.847671  
[AbstractPlus](#) | Full Text: [PDF\(320 KB\)](#) [IEEE CNF](#)  
[Rights and Permissions](#)

**IEEE CNF** IEEE Conference Proceeding**IEE CNF** IEE Conference Proceeding**IEEE STD** IEEE Standard

- 2. **A novel conflict detection and resolution strategy based on TLRSP in replicated database systems**  
Zhiming Ding; Xiaofeng Meng; Shan Wang;  
[Database Systems for Advanced Applications, 2001. Proceedings. Seventh Int'l Conference on](#)  
18-21 April 2001 Page(s):234 - 240  
Digital Object Identifier 10.1109/DASFAA.2001.916383  
[AbstractPlus](#) | Full Text: [PDF\(456 KB\)](#) [IEEE CNF](#)  
[Rights and Permissions](#)

**Indexed by**  
[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -